



UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE
United States Patent and Trademark Office
Address: COMMISSIONER FOR PATENTS
P.O. Box 1450
Alexandria, Virginia 22313-1450
www.uspto.gov

APPLICATION NO.	FILING DATE	FIRST NAMED INVENTOR	ATTORNEY DOCKET NO.	CONFIRMATION NO.
-----------------	-------------	----------------------	---------------------	------------------

10/669,931

09/24/2003

Hans F. van Rietschote

5760-13900

4596

35690

7590

07/17/2007

MEYERTONS, HOOD, KIVLIN, KOWERT & GOETZEL, P.C.

P.O. BOX 398

AUSTIN, TX 78767-0398

EXAMINER

PUENTE, EMERSON C

ART UNIT

PAPER NUMBER

2113

MAIL DATE

DELIVERY MODE

07/17/2007

PAPER

Please find below and/or attached an Office communication concerning this application or proceeding.

The time period for reply, if any, is set in the attached communication.



UNITED STATES PATENT AND TRADEMARK OFFICE

Commissioner for Patents
United States Patent and Trademark Office
P.O. Box 1450
Alexandria, VA 22313-1450
www.uspto.gov

MAILED

JUL 17 2007

Technology Center 2100

**BEFORE THE BOARD OF PATENT APPEALS
AND INTERFERENCES**

Application Number: 10/669,931
Filing Date: September 24, 2003
Appellant(s): VAN RIETSCHOTE ET AL.

Lawrence J. Merkel
Reg. No. #41,191
For Appellant

EXAMINER'S ANSWER

This is in response to the appeal brief filed March 20, 2007 appealing from the Office action mailed October 26, 2006.

(1) Real Party in Interest

A statement identifying by name the real party in interest is contained in the brief.

(2) Related Appeals and Interferences

The examiner is not aware of any related appeals, interferences, or judicial proceedings which will directly affect or be directly affected by or have a bearing on the Board's decision in the pending appeal.

(3) Status of Claims

The statement of the status of claims contained in the brief is correct.

(4) Status of Amendments After Final

The appellant's statement of the status of amendments after final rejection contained in the brief is correct.

(5) Summary of Claimed Subject Matter

The summary of claimed subject matter contained in the brief is correct.

(6) Grounds of Rejection to be Reviewed on Appeal

The appellant's statement of the grounds of rejection to be reviewed on appeal is substantially correct. Appellant indicated in the appeal brief that the grounds of rejection for claim 25 was incorrect (see page 10). Claim 25 should have been rejected over Vert in view of Mashayekhi and Dinker, the third grounds of rejection cited. In a phone interview on July 3, 2007, appellant agreed it was not necessary to reopen prosecution of application as such a change would not require the addition of new prior art or alter the issues at hand. Appellant approved examiner's modification of grounds of rejection to remove claim 25 from the second grounds of

rejection (Vert in view of Mashayekhi) and place claim 25 in the third grounds of rejection (Vert in view of Mashayekhi and Dinker) in order to expedite prosecution of the case to the board.

The changes are as follows:

2. Claims 1-4, 8, 10-13, 19-22, 27-34, and 36-41 are rejected under 35 U.S.C. §103(a) as being unpatentable over Vert et al., U.S. Patent No. 6,360,331 ("Vert") in view of Mashayekhi et al., U.S. Patent No. 6,922,791 ("Mashayekhi").
3. Claims 5-6, 14-15, 18, 23, and 25 are rejected under 35 U.S.C. § 103(a) as being unpatentable over Vert in view of Mashayekhi and Dinker et al., U.S. Patent No. 6,944,788 ("Dinker").

(7) Claims Appendix

The copy of the appealed claims contained in the Appendix to the brief is correct.

(8) Evidence Relied Upon

6,629,266	Harper et al.	9-2003
6,360,331	Vert et al.	3-2002
6,922,791	Mashayekhi et al.	7-2005
6,944,788	Dinker et al.	9-2005

"Collegiate Dictionary", 10th ed, 2001, Merriam-Webster, Inc, pp 938 term "provision".

(9) Grounds of Rejection

As indicated above, Claim 25 has been removed from the rejection Vert in view of Mashayekhi and included in rejection Vert in view of Mashayekhi and Dinker, which was approved in a telephone interview with appellant on July 3, 2007.

The following ground(s) of rejection are applicable to the appealed claims:

Claim Rejections - 35 USC § 102

The following is a quotation of the appropriate paragraphs of 35 U.S.C. 102 that form the basis for the rejections under this section made in this Office action:

A person shall be entitled to a patent unless –

(e) the invention was described in (1) an application for patent, published under section 122(b), by another filed in the United States before the invention by the applicant for patent or (2) a patent granted on an application for patent by another filed in the United States before the invention by the applicant for patent, except that an international application filed under the treaty defined in section 351(a) shall have the effects for purposes of this subsection of an application filed in the United States only if the international application designated the United States and was published under Article 21(2) of such treaty in the English language.

Claims 16 and 17 are rejected under 35 U.S.C. 102(e) as being anticipated by US Patent No. 6,629,226 of Harper et al. referred hereinafter “Harper”.

In regards to claim 16, Harper discloses:

detecting that an application in a first node is to failover (see figure 5 and column 8 lines 11-16);

provisioning a second node to execute the application responsive to the detecting (see figure 5 and column 8 lines 11-16);

attempting to failing the application over from the first node to the second node (see figure 5 item 502 and column 7 lines 62-67).

detecting a lack of success in the failover, wherein the lack of success is due to a lack of an eligible node (see figure 5 item 502 and column 7 lines 62-67); and

permitting the application to execute on the first node responsive to the lack of the eligible node if the attempt to failover is due to a performance of the application on the first node being less than a threshold performance level. Harper disclose restarting application on a second node if a fail-to node is available (see figure 5 and column 8 lines 11-16), implying that the

Art Unit: 2113

application is still run on the first node if a second node is not available. Harper further discloses the rejuvenation is done when parameters approach an exhaustion threshold level (see column 8 lines 60-61).

In regards to claim 17, Harper discloses:

wherein, if the attempt to failover is due to a failure in a service group including the application, the method further comprises notifying an administrator. Harper discloses trigger failover when one or more parameters reaches a hazardous region (see column 8 lines 60-61) and it is determined that a fail-to node is unavailable, notifying an operator (see column 8 lines 1-3).

Claim Rejections - 35 USC § 103

The following is a quotation of 35 U.S.C. 103(a) which forms the basis for all obviousness rejections set forth in this Office action:

(a) A patent may not be obtained though the invention is not identically disclosed or described as set forth in section 102 of this title, if the differences between the subject matter sought to be patented and the prior art are such that the subject matter as a whole would have been obvious at the time the invention was made to a person having ordinary skill in the art to which said subject matter pertains. Patentability shall not be negated by the manner in which the invention was made.

Claims 1-4, 8,10-13,19-22, 27-34, and 36-41 are rejected under 35 U.S.C. 103(a) as being unpatentable over US Patent No. 6,360,331 of Vert et al. referred hereinafter "Vert" in view of US Patent No. 6,922,791 of Mashayekhi et al. referred hereinafter "Mashayekhi".

In regards to claim 1 and 19, Vert discloses:

detecting that an application in a first node is to failover, wherein the first node is included in a cluster being used to execute the application (see column 2 lines 35-41);

However, Vert fails to explicitly disclose:

adding a second node to the cluster responsive to the detecting, provisioning the second node to execute the application responsive to the detecting and failing the application over from the first node to the second node.

Mashayekhi discloses a commonly known and used failover policy wherein a passive node provides failover for all active nodes in the cluster (see column 2 lines 60-67), indicating adding a second node to the cluster responsive to the detecting, provisioning the second node to execute the application responsive to the detecting and failing the application over from the first node to the second node. Examiner interprets “the cluster” described in the claims as a cluster of active node(s). Since a passive node is initially tasked with running no applications, a passive node is not functioning as part of a cluster.

It would have been obvious to one of ordinary skill in the art at the time the invention was made to modify the teachings of Vert with that of Mashayekhi wherein the second node is a passive node that provides failover for all active nodes in the cluster (see column 2 lines 60-67), indicating adding a second node to the cluster responsive to the detecting, provisioning the second node to execute the application responsive to the detecting and failing the application over from the first node to the second node. A person of ordinary skill in the art would have been motivated to combine the teachings of Vert and Mashayekhi because Vert is concerned with providing failover (see column 5 lines 48-52) and a passive node that provides failover for all active nodes in the cluster, as per teachings of Mashayekhi, constitutes a commonly known and used failover policy (see column 2 lines 60-67). Furthermore, such a failover policy would be advantageous since it would not provide additional workload to active nodes upon failure.

In regards to claim 2 and 20, Vert discloses:

activating one or more resources used by the application on the second node (see column 2 lines 35-41 and column 7 lines 35-40).

In regards to claim 3 and 21, Vert discloses:

wherein the provisioning comprises installing one or more resources used by the application on the second node (see column 2 lines 35-41 and column 7 lines 35-40).

In regards to claim 4 and 22, Vert discloses:

wherein the second node has multiple boot capability (see page 9 lines 30-31), and wherein the provisioning comprises rebooting the second node from a partition that comprises one or more resources used by the application (see column 9 lines 5-11).

In regards to claim 8, Vert discloses:

adding the first node to the plurality of nodes to be selectable for provisioning (see column 4 line 63 to column 5 line 5 and column 9 lines 29-31).

In regards to claim 10 and 27, Vert discloses:

wherein the detecting comprises detecting that the performance of the application executing on the first node is less than a threshold performance level. Vert discloses sending periodic messages, called heartbeats, to detect the communication path is good and other system are operational (see column 5 lines 30-35). In the event of a communication failure (no heartbeat), the system fails over to one or more active systems (see column 5 lines 48-52). The heartbeats represent the performance of the application. When heartbeats are not received, the application performance on the first node is less than threshold performance level.

In regards to claim 11 and 28, Vert discloses:

wherein the performance is less than the threshold performance level for at least a predetermined time interval. Vert discloses sending periodic messages, called heartbeats, to detect the communication path is good and other system are operational (see column 5 lines 30-35). In the event of a communication failure (no heartbeat), the system fails over to one or more active systems (see column 5 lines 48-52). The heartbeats represent the performance of the application. When heartbeats are not received within a period of time, the application performance on the first node is less than threshold performance level for at least a predetermined time interval.

In regards to claim 12 and 29, Vert discloses:

wherein the detecting comprises alternatively detecting a failure in a service group including the application (see column 9 lines 5-15).

In regards to claim 13 and 30, Vert discloses:

wherein the detecting comprises detecting a failure in a service group including the application (see column 9 lines 5-15).

In regards to claim 31, Vert discloses:

a plurality of nodes, wherein a first node of the plurality of nodes is configured to monitor performance of an application executing on a second node of the plurality of nodes during use and wherein the second node is included in a cluster being used to execute the application (see column 5 lines 40-45), and

However, Vert fails to explicitly disclose:

wherein, in response to a detection that the application is to failover from the first node, a third node is configured to be provisioned to execute the application, wherein the third node is

added to the cluster responsive to detecting that the application is to failover from the second node during use, and wherein the application is failed over to the third node during use.

Mashayekhi discloses a commonly known and used failover policy wherein a passive node provides failover for all active nodes in the cluster (see column 2 lines 60-67), indicating in response to a detection that the application is to failover from the first node, a third node is configured to be provisioned to execute the application, wherein the third node is added to the cluster responsive to detecting that the application is to failover from the second node during use, and wherein the application is failed over to the third node during use. Examiner interprets “the cluster” described in the claims as a cluster of active node(s). Since a passive node is initially tasked with running no applications, a passive node is not functioning as part of a cluster.

It would have been obvious to one of ordinary skill in the art at the time the invention was made to modify the teachings of Vert with that of Mashayekhi wherein the second node is a passive node that provides failover for all active nodes in the cluster (see column 2 lines 60-67), indicating in response to a detection that the application is to failover from the first node, a third node is configured to be provisioned to execute the application, wherein the third node is added to the cluster responsive to detecting that the application is to failover from the second node during use, and wherein the application is failed over to the third node during use. A person of ordinary skill in the art would have been motivated to combine the teachings of Vert and Mashayekhi because Vert is concerned with providing failover (see column 5 lines 48-52) and a passive node that provides failover for all active nodes in the cluster, as per teachings of Mashayekhi, constitutes a commonly known and used failover policy (see column 2 lines 60-67).

Art Unit: 2113

Furthermore, such a failover policy would be advantageous since it would not provide additional workload to active nodes upon failure.

In regards to claim 32, Vert discloses:

activating one or more resources used by the application on the second node (see column 2 lines 35-41 and column 7 lines 35-40).

In regards to claim 33, Vert discloses:

wherein the provisioning comprises installing one or more resources used by the application on the second node (see column 2 lines 35-41 and column 7 lines 35-40).

In regards to claim 34, Vert discloses:

wherein the second node has multiple boot capability (see page 9 lines 30-31), and wherein the provisioning comprises rebooting the second node from a partition that comprises one or more resources used by the application (see column 9 lines 5-11).

In regards to claim 36, Vert discloses:

wherein the first node is configured to detect that the performance of the application executing on the second node is less than a threshold performance level. Vert discloses sending periodic messages, called heartbeats, to detect the communication path is good and other system are operational (see column 5 lines 30-35). In the event of a communication failure (no heartbeat), the system fails over to one or more active systems (see column 5 lines 48-52). The heartbeats represent the performance of the application. When heartbeats are not received, the application performance on the first node is less than threshold performance level.

In regards to claim 37, Vert discloses:

Art Unit: 2113

wherein the performance is less than the threshold performance level for at least a predetined time interval. Vert discloses sending periodic messages, called heartbeats, to detect the communication path is good and other system are operational (see column 5 lines 30-35). In the event of a communication failure (no heartbeat), the system fails over to one or more active systems (see column 5 lines 48-52). The heartbeats represent the performance of the application. When heartbeats are not received within a period of time, the application performance on the first node is less than threshold performance level for at least a predetermined time interval.

In regards to claim 38, Vert discloses:

wherein the second node is configured to detect a failure in a service group including the application, and wherein the application is to failover from the second node if the second node detects the failure (see column 9 lines 5-15).

In regards to claim 39 and 40, Vert discloses:

removing the first node from the cluster responsive to successfully failing over the application to the second node (see column 5 lines 48-49).

In regards to claim 41, Vert discloses:

wherein the second node is removed from the cluster responsive to successful failover to the third node (see column 5 lines 48-49).

Claims 5, 6, 14, 15, 18, 23, and 25 are rejected under 35 U.S.C. § **103(a)** as being unpatentable over Vert in view of Mashayekhi and in further view of US Patent No. 6,944,788 of Dinker et al. referred hereinafter "Dinker".

In regards to claim 5 and 23, Vert in view of Mashayekhi fails to explicitly disclose:

Art Unit: 2113

selecting the second node from a plurality of nodes;

However, Dinker discloses a plurality of backup and alternate application servers or nodes for failover, indicating selecting the second node from a plurality of nodes (see column 9 lines 10-16).

It would have been obvious to one of ordinary skill in the art at the time the invention was made to combine the teaching of Vert, Mashayekhi, and Dinker to have a plurality of nodes for failover, indicating selecting the second node from a plurality of nodes. A person of ordinary skill in the art would have been motivated to combine the teachings because Vert is concerned with providing failover (see column 2 lines 34-41) and having a plurality of nodes for failover, as per teachings of Dinker, provides additional levels of failover (see column 9 lines 10-16).

In regards to claim 6, Mashayekhi discloses:

wherein the second node is executing a different application when selected (see column 9 lines 22-27). Mashayekhi discloses wherein the cluster is not running any applications prior to failover and running applications actively when there is a failover (see column 2 lines 65-67).

In regards to claim 14, Vert in view of Mashayekhi fails to explicitly disclose:

detecting a lack of success in the failing over.

However, Dinker discloses detecting a lack of success in the failing over. Dinker discloses a primary application server and one or more backup application servers (see column 8 lines 30-33). Dinker further discloses when the primary becomes unavailable, a first backup is promoted to the role of the new primary (see column 8 lines 30-33). Thus, when the new primary fails or becomes unavailable, a second backup becomes the new primary. The instance

Art Unit: 2113

when the first backup that becomes the new primary fails constitutes a lack of success in the failing over.

It would have been obvious to one of ordinary skill in the art at the time the invention was made to combine the teaching of Vert, Mashayekhi, and Dinker to detect a lack of success in the failing over. A person of ordinary skill in the art would have been motivated to combine the teachings because Vert is concerned with providing failover (see column 2 lines 34-41) and providing more backups, as per teachings of Dinker, provides additional levels of failover (see column 8 lines 30-33).

In regards to claim 15, Dinker discloses:

provisioning a third node to execute the application responsive to detecting the lack of success, and failing the application over from the second node to the third node. Dinker discloses a primary application server and one or more backup application server (see column 8 lines 30-33). Dinker further discloses when the primary becomes unavailable, one of the backups is promoted to the role of the new primary (see column 8 lines 30-33). When the new primary fails or becomes unavailable, another backup becomes the new primary, indicating provisioning a third node to execute the application responsive to detecting the lack of success, and failing the application over from the second node to the third node

In regards to claim 18, Vert in view of Mashayekhi fails to explicitly disclose:

determining that a performance level on the second node is less than a threshold;
provisioning a third node to execute the application responsive to the determining;
failing the application over from the second node to the third node.

Art Unit: 2113

Dinker discloses a primary application server and one or more backup application servers (see column 8 lines 30-33). Dinker further discloses when the primary becomes unavailable, one of the backups is promoted to the role of the new primary (see column 8 lines 30-33). When the new primary fails or becomes unavailable, another backup becomes the new primary, indicating provisioning a third node to execute the application responsive to the determining, and failing the application over from the second node to the third node. Dinker also discloses the fact the primary application is unreachable may be discovered by a heartbeat mechanism (see column 11 lines 49-51). When heartbeats are not received, the application's performance level on the second node is less than a threshold.

It would have been obvious to one of ordinary skill in the art at the time the invention was made to combine the teaching of Vert, Mashayekhi, and Dinker to determine that a performance level on the second node is less than a threshold, provision a third node to execute the application responsive to the determining, and failing the application over from the second node to the third node. A person of ordinary skill in the art would have been motivated to combine the teachings because Vert is concerned with failover (see column 2 lines 34-41) and providing more backups (see column 8 lines 30-33), as per teachings of Dinker, provides additional levels of failover.

In regards to claim 25, Vert discloses:

adding the first node to the plurality of nodes to be selectable for provisioning (see column 4 line 63 to column 5 line 5 and column 9 lines 29-31).

Claims 7 and 24 are rejected under 35 U.S.C. § 103(a) as being unpatentable over Vert in view of Mashayekhi and Dinker and in further view of Harper.

In regards to claim 7 and 24, Vert in view of Mashayekhi and Dinker fails to explicitly disclose:

wherein the selecting comprises verifying that the second node includes hardware that is sufficient to execute the application.

However, Harper discloses wherein the selecting comprises verifying that the second node includes hardware that is sufficient to execute the application (see figure 5 and column 7 lines 60-67).

It would have been obvious to one of ordinary skill in the art at the time the invention was made to combine the teachings of Vert, Mashayekhi, and Harper to verifying that the second node includes hardware that is sufficient to execute the application. A person of ordinary skill in the art would have been motivated to combine the teachings because is Vert is concerned with providing failover (see column 5 lines 48-52) and verifying that the second node includes hardware that is sufficient to execute the application, as per teachings of Harper, ensures the second node is capable of providing failover (see figure 5 and column 7 lines 60-67).

(10) Response to Argument

First Grounds of Rejection:

Claims 16-17

In response to appellant's argument, "These teachings relate to failing over the node to an (already provisioned) secondary node. Accordingly, these teachings have nothing to do with the provisioning features recited in claim 16.

Harper teaches that the primary node and the secondary (or backup) node are both configured to execute the application when the cluster is created. See, e.g., col. 6, lines 32-42: “Typically, in a two-node cluster, one node is designated the 'primary node' and normally runs the application software, and another is designated the 'backup node' and is capable of running the application when the primary node fails. Distributed cluster management software running on both the primary node and the secondary node continually checks on the health of the primary node and its associated application software.” Therefore, Harper teaches that the backup/secondary node is provisioned to execute the application before execution begins on the primary node.

Nothing in Harper teaches or suggests “detecting that an application in a first node is to failover; [and] provisioning a second node to execute the application responsive to the detecting” as recited in claim 16.” (see page 5-6) examiner respectfully disagrees.

Harper disclose when there is a failure to the primary and the fail-to node can accept a failover workload, the rejuvenation agent on the primary node instructs the cluster manager to gracefully shut down and restart the application on the secondary computer (see column 8 lines 11-16). Examiner interprets the failover process to enable restarting of the application on the secondary computer as “provisioning a second node to execute the application responsive to the detecting”. The term “provisioning” is defined in Webster’s Dictionary as “the act or process of providing” (see page 938 term “provision”). As Harper discloses providing the secondary computer to run the application upon a failure to the primary, Harper discloses “the act or process of providing”, indicating provisioning. Argument is moot. Examiner maintains his rejection.

Second Grounds of Rejection:

Claims 1-3,13,19-21,30-33, and 38-41

In response to appellant's argument, "The Office Action alleges that Mashayekhi teaches the above highlighted features at col. 2, lines 60-67, asserting that the Examiner interprets "cluster" in the claims to be a cluster of active nodes as described in Mashayekhi. Appellant respectfully disagrees. The interpretation asserted by the Office Action is clearly contradicted by the plain language in Mashayekhi, in which the passive node is clearly part of the cluster and there is no cluster of active nodes the excludes the passive node. For example, Mashayekhi teaches: "Another known failover policy utilizes a separate 'passive' node that is present in the cluster exclusively for the purpose of being the failover node for all active nodes in the cluster. As illustrated in the following graph, each node on the cluster that is actively running applications (nodes 1-3) fails over to node 4, which is not tasked with running any applications other than in the event of a failover." (Mashayekhi, col. 2, lines 60-67). Thus, it is clear that Mashayekhi's cluster is four nodes, three of which are active and one of which is passive. All four nodes are clearly part of the cluster, and the passive node is provisioned a priori to execute any application from nodes 1 to 3 in the event of a failover. Thus, in the cited section, all that occurs when a failover event is detected is the act of failing over itself.

Accordingly, the cited section of Mashayekhi does not teach or suggest "detecting that an application in a first node is to failover, wherein the first node is included in a cluster being used to execute the application; adding a second node to the cluster responsive to the detecting; [and] provisioning the second node to execute the application responsive to the detecting " as recited in claim 1. Vert does not teach or suggest the above highlighted features, either. Accordingly, the

alleged combination of Vert and Mashayekhi does not teach or suggest the combination of features recited in claim 1” (see page 7-8) examiner respectfully disagrees.

As seen in Fig 2 and page 5 top paragraph of appellant’s Specification, the clusters include a node or nodes executing applications in a steady state, clearly defining the cluster as a “cluster of active node(s)”. Likewise, Mashayekhi discloses a plurality of nodes, each running an application, thus indicating a “cluster of active nodes” (see column 2 lines 64-66). Mashayekhi further discloses a passive node, which is not tasked with running any applications other than in the event of the failover (see column 2 lines 60-61 and 66-67). As the passive node is not running any applications prior to failure, the passive node is not part of the cluster or cluster of active nodes. When there is a failure to one of the active nodes, the passive node becomes active and is tasked to run the application of the failed active node, thus becoming part of the cluster or cluster of active nodes, indicating adding a second node to the cluster responsive to the detecting. As indicated above, the term “provisioning” is defined in Webster’s Dictionary as “the act or process of providing” (see page 938 term “provision”). As Mashayekhi discloses providing the secondary computer to run the application upon a failure to the primary (see column 2 lines 60-61 and 66-67), Vert in view of Mashayeki discloses “the act or process of providing”, indicating provisioning the second node to execute the application responsive to the detecting. Argument is moot. Examiner maintains his rejection.

The arguments pertaining to independent claim 19 and 31 and dependent claims 2-3,13,20-21,30,32-33, and 38-41 on page 8 are similar to that of claim 1 and as such, are rejected for the same reasons.

Claims 4, 22, and 34

In response to appellant's argument, "Vert fails to teach or suggest 'the second node has multiple boot capability, and wherein the provisioning comprises rebooting the second node from a partition that comprises one or more resources used by the application'," (see page 9) examiner respectfully disagrees.

Vert discloses being able to take resources or applications offline and restarting (rebooting) groups on another system (partition), indicating rebooting the second node from a partition that comprises one or more resources used by the application. Furthermore, Vert discloses being able to bring previous failed systems back online (see page 9 lines 30-31), implying the restarting or rebooting of systems, indicating a multiple boot capacity.

Claims 8 and 25

In response to appellant's argument that Vert fails to teach "adding the first node (from which the application has failed over) to the plurality of nodes that can be provisioned and added to the cluster" (see pages 10-11) examiner respectfully disagrees.

Vert discloses during creation of a new cluster, adding members to the cluster (see column 4 lines 63-67). It is understood the first node must be added initially in order to function as part of the cluster, indicating adding the first node to the nodes that can be provisioned and added to another cluster. Examiner notes claims do not include limitation "from which the application has failed over". However, even if claims included such limitation, Vert further discloses being able to bring previous failed systems back online (see page 9 lines 30-31), indicating adding the first node (from which the application has failed over) to the nodes that can

Art Unit: 2113

be provisioned and added to another cluster. Argument is moot. Examiner maintains his rejection.

Claims 10,12,27,29, and 36

In response to appellant's argument that Vert fails to teach, "detecting that the performance of the application executing on the first node is less than a threshold performance level," (see pages 11-12) examiner respectfully disagrees.

Vert discloses sending periodic messages, called heartbeats, to detect the communication path is good and other system are operational (see column 5 lines 30-35). In the event of a communication failure (no heartbeat), the system fails over to one or more active systems (see column 5 lines 48-52). The heartbeats represent the performance of the application. When heartbeats are not received, the application performance on the first node is less than threshold performance level. Argument is moot. Examiner maintains his rejection.

Claims 11,28, and 37

In response to appellant's argument that Vert fails to teach, "the performance is less than the threshold performance level for at least a predefined time interval," (see pages 13) examiner respectfully disagrees.

Vert discloses sending periodic messages, called heartbeats, to detect the communication path is good and other system are operational (see column 5 lines 30-35). In the event of a communication failure (no heartbeat), the system fails over to one or more active systems (see column 5 lines 48-52). The heartbeats represent the performance of the application. When heartbeats are not received within a period of time, the application performance on the first node

is less than threshold performance level for at least a predetermined time interval. Argument is moot. Examiner maintains his rejection.

Third Grounds of Rejection:

Claims 5, 14, and 23

The arguments pertaining to independent claim 5, 14, and 23 on page 14 are similar to that of base claims 1 and 19 and as such, are rejected for the same reasons.

Claim 6

In response to appellant's argument that prior art fails to teach, "the second node is executing a different application when selected [to be provisioned for failover of the application from the first node]," (see pages 14) examiner respectfully disagrees.

Mashayekhi discloses wherein the node is running no applications prior to failover and running applications actively when there is a failover (see column 2 lines 65-67), indicating the second node is executing a different application when selected. Argument is moot. Examiner maintains his rejection.

Claim 15

In response to appellant's argument that prior art fails to teach, "provisioning a third node to execute the application responsive to detecting the lack of success, and failing the application over from the second node to the third node," (see pages 15) examiner respectfully disagrees.

As indicating above, Vert in view of Mashayekhi discloses provisioning in response to detecting a failover. Dinker further discloses when the primary becomes unavailable, one of the backups is promoted the role of the new primary (see column 8 lines 30-33). When the new primary fails or becomes unavailable, another backup becomes the new primary, indicating

provisioning a third node to execute the application responsive to detecting the lack of success, and failing the application over from the second node to the third node. Argument is moot.

Examiner maintains his rejection.

Claim 18

In response to appellant's argument that prior art fails to teach, "determining that a performance level on the second node is less than a threshold; provisioning a third node to execute the application responsive to determining, and failing the application over from the second node to the third node," (see pages 16) examiner respectfully disagrees.

As indicated above, Vert in view of Mashayekhi discloses provisioning in response to detecting a failover. Vert also discloses sending periodic messages, called heartbeats, to detect the communication path is good and other system are operational (see column 5 lines 30-35). In the event of a communication failure (no heartbeat), the system fails over to one or more active systems (see column 5 lines 48-52). The heartbeats represent the performance of the application. When heartbeats are not received, the application performance on the first node is less than threshold performance level. Dinker further discloses when the primary becomes unavailable, one of the backups is promoted to the role of the new primary (see column 8 lines 30-33). When the new primary fails or becomes unavailable, another backup becomes the new primary, indicating determining that a performance level on the second node is less than a threshold and provisioning a third node to execute the application responsive to determining, and failing the application over from the second node to the third node. Argument is moot. Examiner maintains his rejection.

Fourth Grounds of Rejection:

Art Unit: 2113

Claims 7 and 24

The arguments pertaining to independent claim 7 and 24 on page 17 are similar to that of base claims 1 and 19 and as such, are rejected for the same reasons.

(11) Related Proceeding(s) Appendix

No decision rendered by a court or the Board is identified by the examiner in the Related Appeals and Interferences section of this examiner's answer.

For the above reasons, it is believed that the rejections should be sustained.

Respectfully submitted,



Emerson Puente
Examiner
AU 2113

Conferees:

Robert Beausoliel



 MARK H. RINEHART
SUPERVISORY PATENT EXAMINER
TECHNOLOGY CENTER 2100